THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Optimal Control of Multi-phase Movements with Learned Dynamics

Citation for published version:

Radulescu, A, Nakanishi, J & Vijayakumar, S 2016, Optimal Control of Multi-phase Movements with Learned Dynamics. in A Gruca, A Brachman, S Kozielski & T Czachórski (eds), *Man–Machine Interactions 4: 4th International Conference on Man–Machine Interactions, ICMMI 2015 Kocierz Pass, Poland, October 6–9, 2015*. Advances in Intelligent Systems and Computing, vol. 391, Springer International Publishing, Cham, pp. 61–76, 4th International Conference on Man–Machine Interactions, Kocierz Pass, Poland, 6/10/15. https://doi.org/10.1007/978-3-319-23437-3_5

Digital Object Identifier (DOI):

[10.1007/978-3-319-23437-3_5](https://doi.org/10.1007/978-3-319-23437-3_5)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Man–Machine Interactions 4

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Optimal Control of Multi-Phase Movements with Learned Dynamics

Andreea Radulescu¹, Jun Nakanishi², Sethu Vijayakumar¹

¹ Institute of Perception, Action and Behaviour, University of Edinburgh, UK

² Institute for Cognitive Systems, Technical University of Munich, Germany

Abstract. In this paper, we extend our work on movement optimisation for variable stiffness actuation (VSA) with multiple phases and switching dynamics to incorporate scenarios with incomplete, complex or hard to model robot dynamics. By incorporating a locally weighted nonparametric learning method to model the discrepancies in the system dynamics, we formulate an online adaptation scheme capable of systematically improving the multi-phase plans (stiffness modulation and torques) and switching instances while improving the dynamics model on the fly. This is demonstrated on a realistic model of a VSA brachiating system with excellent adaptation results.

1 Introduction

The accuracy of model-based control is significantly dependent on that of the models themselves. Traditional robotics employs models obtained from mechanical engineering insights. Kinematic equations will provide accurate information about the evolution of a rigid body configuration, given a precise knowledge of its geometry. Similarly, the dynamics equation can incorporate well modelled factors such as inertia, Coriolis and centrifugal effect or external forces.

However, there are certain elements that cannot be fully captured by these models, such as friction from the joints or resulting from cable movement [1], which can vary in time. The introduction of flexible elements in the structure of a system increases the complexity of the model and makes the identification of accurate dynamics significantly more difficult. Additionally, during operation, the robot can suffer changes in its mechanical structure due to wear and tear or due to the use of a tool (which modifies the mechanical chain structure) [2].

On-line adaptation of models can provide a solution for capturing all these properties. Early approaches, such as on-line *parameter identification* [3], which tunes the parameters of a predefined model (dictated by the mechanical structure) using data collected during operation, proved sufficient for accurate control and remained a popular approach for a long time [4, 5]. The increased complexity of latest robotic systems demands novel approaches capable of accommodating significant non-linear and unmodelled robot dynamics. Successful *non-parametric model learning* methods use supervised learning to perform system

identification with only limited prior information about its structure - removing the restriction to a fixed model structure, allowing the model complexity to adapt in a data driven manner.

In this work, we will build on our significant prior efforts to engage this techniques in the context of robot control [6, 7, 8] and apply this in the context of multiphase variable impedance movements. Indeed, adaptive model learning has been used successfully in a wide range of scenarios such as inverse dynamics control [9, 10], inverse kinematics [11, 12], robot manipulation and locomotion [13, 14].

1.1 Adaptive Learning for Optimal Control

Classical OC (Optimal Control) is formulated using an analytic dynamics model, but recent work [15, 16] has shown that combining OC with dynamics learning can produce a powerful and principled control strategy for complex systems with redundant actuation.

In [16], using online (non-parametric) supervised learning methods, an adaptive internal model of the system dynamics is learned. The model is afterwards used to derive an optimal control law. This approach, named *iterative Linear Quadratic-Gaussian (iLQG) method with learned dynamics (iLQG-LD)*, proved efficient in a variety of realistic scenarios, including problems where the analytic dynamics model is difficult to estimate accurately or subject to changes and the system is affected by noise [16, 17].

In iLQG-LD the update of the dynamics model takes place on a trial-by-trial basis [17]. The operating principle (depicted in Fig. 1) is to (i) compute the iLQG solution, (ii) run the obtained control law on the plant and collect data, (iii) use the plant data to update the dynamics model.

Fig. 1. The iLQG-LD learning and control scheme as first introduced in [17]

The initial state and the cost function (which includes the desired final state) are provided to the iLQG planner, alongside a preliminary model of the dynamics. An initial (locally optimal) command sequence $\bar{u}$ is generated, together with the corresponding state sequence $\bar{x}$ and feedback correction gains $\mathbf{L}$. Applying the feedback controller scheme, at each time step the control command is corrected by $\delta \mathbf{u} = \mathbf{L}(\mathbf{x} - \bar{\mathbf{x}})$, where $\mathbf{x}$ is the true state of the plant. The model of the

dynamics is updated using the information provided by the applied command $\mathbf{u} + \delta\mathbf{u}$ and observed state $\mathbf{x}$.

This methodology employs the *Locally Weighted Projection Regression (LWPR)* [18] as the nonparametric learning scheme of choice to train a model of the dynamics in an incremental fashion. In LWPR, the regression function is constructed by combining local linear models. During training the parameters of the local models (locality and fit) are updated using incremental *partial least squares (PLS)*. PLS projects the input on to a small number of directions in the input space along the directions of maximal correlation with the output and then performs linear regression on the projected inputs. This makes LWPR suitable for high dimensional input spaces. Local models can be pruned or added on an as-need basis (e.g., when training data is generated in previously unexplored regions). The areas of validity (receptive fields) of each local model are modelled by Gaussian kernels. LWPR keeps a number of variables that hold sufficient statistics for the algorithm to perform the required calculations incrementally.

We incorporate the iLQG-LD scheme into our approach involving learning the dynamics of a brachiation system with VSA (variable stiffness actuator) capabilities and employing it in planning for locomotion tasks.

2 Problem Formulation

In our previous work [19], we introduced a general formulation of optimal control problems for tasks with multiple phase movements including switching dynamics and discrete state transition arising from interactions with an environment. Given a rigid body dynamics formulation of a robot with a VSA model, a hybrid dynamics representation with a composite cost function is introduced to describe such a task. In this section we briefly describe this approach, for details we refer the interested reader to [19]. We also introduce the changes dictated by the use of the LWPR method in the context of iLQG-LD, for integration within our approach.

2.1 Hybrid Dynamics with Time-based Switching and Discrete State Transition

We employ the following hybrid dynamics representation to model multi-phase movements having interactions with an environment [20]:

$$\dot{\mathbf{x}} = \mathbf{f}_{i_j}(\mathbf{x}, \mathbf{u}), \quad T_j \leq t < T_{j+1} \quad (1)$$

$$\mathbf{x}(T_j^+) = \Delta^{i_{j-1}, i_j}(\mathbf{x}(T_j^-)) \quad (2)$$

with $j = 0, \dots, K$ for (1) and $j = 1, \dots, K$ for (2) and where $\mathbf{f}_i : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the i -th subsystem, $\mathbf{x} \in \mathbb{R}^n$ is a state vector, $\mathbf{u} \in \mathbb{R}^m$ is a control input vector.

When the dynamics switch from subsystem i_{j-1} to i_j at $t = T_j$, we assume that instantaneous discrete (discontinuous) state transition is introduced, which is denoted by a map Δ^{i_{j-1}, i_j} in (2). The terms $\mathbf{x}(T_j^+)$ and $\mathbf{x}(T_j^-)$ denote the post-

and pre-transition states, respectively. In this case, the sequence of switching is assumed to be given. Fig. 2 depicts a schematic diagram of a hybrid system we consider in this work.

Fig. 2. A hybrid system with time-based switching dynamics and discrete state transition with a known sequence. The objective is to find an optimal control command $\mathbf{u}$, switching instances T_i and final time T_f which minimises the composite cost J .

2.2 Robot Dynamics with Variable Stiffness Actuation

To describe multi-phase movements, we consider multiple sets of robot dynamics, as described by (1). An individual rigid body dynamics model is defined for each associated phase of the movement as a subsystem. The servo motor dynamics in the VSA are modelled as a critically damped second order dynamics:

$$\mathbf{M}_i(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}_i(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}_i(\mathbf{q}) + \mathbf{D}_i\dot{\mathbf{q}} = \boldsymbol{\tau}_i(\mathbf{q}, \mathbf{q}_m) \quad (3)$$

$$\ddot{\mathbf{q}}_m + 2\boldsymbol{\alpha}_i\dot{\mathbf{q}}_m + \boldsymbol{\alpha}_i^2\mathbf{q}_m = \boldsymbol{\alpha}_i^2\mathbf{u} \quad (4)$$

where i denotes the i -th subsystem, $\mathbf{q} \in \mathbb{R}^n$ is the joint angle vector, $\mathbf{q}_m \in \mathbb{R}^m$ is the motor position vector of the VSA, $\mathbf{M} \in \mathbb{R}^{n \times n}$ is the inertia matrix, $\mathbf{C} \in \mathbb{R}^{n \times n}$ is the Coriolis term, $\mathbf{g} \in \mathbb{R}^n$ is the gravity vector, $\mathbf{D} \in \mathbb{R}^{n \times n}$ is the viscous damping matrix, and $\boldsymbol{\tau} \in \mathbb{R}^n$ are the joint torques from the variable stiffness mechanism. In the equations above, (3) denotes the rigid body dynamics of the robot and (4) denotes the servo motor dynamics in the variable stiffness actuator. In (4), $\boldsymbol{\alpha}$ determines the bandwidth of the servo motors³ and $\mathbf{u}$ is the motor position command [21]. We assume that the range of control command $\mathbf{u}$ is limited between $\mathbf{u}_{min}$ and $\mathbf{u}_{max}$.

In this work, we consider a VSA model in which the joint torques are given in the form

$$\boldsymbol{\tau}(\mathbf{q}, \mathbf{q}_m) = \mathbf{A}^T(\mathbf{q}, \mathbf{q}_m)\mathbf{F}(\mathbf{q}, \mathbf{q}_m) \quad (5)$$

³ $\boldsymbol{\alpha} = \text{diag}\{a_1, \dots, a_m\}$ and $\boldsymbol{\alpha}^2 = \text{diag}\{a_1^2, \dots, a_m^2\}$ for notational convenience.

where $\mathbf{A}$ is the moment arm matrix and $\mathbf{F}$ is the forces by the elastic elements [21] and the joint stiffness is defined as $\mathbf{K} = -\frac{\partial \boldsymbol{\tau}}{\partial \mathbf{q}}$.

We consider the state space representation as the combined plant dynamics consisting of the rigid body dynamics (3) and the servo motor dynamics (4):

$$\dot{\mathbf{x}} = \mathbf{f}_i(\mathbf{x}, \mathbf{u}) \quad (6)$$

where

$$\mathbf{f}_i = \begin{bmatrix} \mathbf{x}_2 \\ \mathbf{M}_i^{-1}(\mathbf{x}_1)(-\mathbf{C}_i(\mathbf{x}_1, \mathbf{x}_2)\mathbf{x}_2 - \mathbf{g}_i(\mathbf{x}_1) - \mathbf{D}_i\dot{\mathbf{x}}_2 + \boldsymbol{\tau}_i(\mathbf{x}_1, \mathbf{x}_3)) \\ \mathbf{x}_4 \\ -\boldsymbol{\alpha}_i^2\mathbf{x}_3 - 2\boldsymbol{\alpha}_i\mathbf{x}_4 + \boldsymbol{\alpha}_i^2\mathbf{u} \end{bmatrix} \quad (7)$$

and $\mathbf{x} = [\mathbf{x}_1^T, \mathbf{x}_2^T, \mathbf{x}_3^T, \mathbf{x}_4^T]^T = [\mathbf{q}^T, \dot{\mathbf{q}}^T, \mathbf{q}_m^T, \dot{\mathbf{q}}_m^T]^T \in \mathbb{R}^{2(n+m)}$ is the state vector consisting of the robot state and the servo motor state.

Employing the iLQG-LD framework we aim to create an accurate model of the dynamics model of the real hardware using supervised learning. We assume the existence of a preliminary analytic dynamics model which takes the form presented in (3, 4), which is inaccurate (due to various factors such as: the inability of the rigid body dynamics to incorporate all the elements of the system's behaviour or changes suffered during operation).

We use the LWPR method to model the error between the true behaviour of the system and the initial model provided. Thus we replace the dynamics $\mathbf{f}_i$ in (6) with the composite dynamics model $\mathbf{f}_{ci}$:

$$\dot{\mathbf{x}} = \mathbf{f}_{ci}(\mathbf{x}, \mathbf{u}) = \tilde{\mathbf{f}}_i(\mathbf{x}, \mathbf{u}) + \bar{\mathbf{f}}_i(\mathbf{x}, \mathbf{u}) \quad (8)$$

where $\tilde{\mathbf{f}}_i \in \mathbb{R}^{2(n+m)}$ is the initial inaccurate model and $\bar{\mathbf{f}}_i \in \mathbb{R}^{2(n+m)}$ is the LWPR model mapping the discrepancy between $\tilde{\mathbf{f}}_i \in \mathbb{R}^{2(n+m)}$ and the behaviour of the system. We note that the changes introduced by iLQG-LD only affect the dynamics modelling in (1), while the instantaneous state transition mapped by Δ in (2) remains unchanged.

2.3 Movement Optimisation of Multiple Phases

For the given hybrid dynamics, in order to describe the full movement with multiple phases, we consider the following composite cost function:

$$J = \phi(\mathbf{x}(T_f)) + \sum_{j=1}^K \psi^j(\mathbf{x}(T_j^-)) + \int_{T_0}^{T_f} h(\mathbf{x}, \mathbf{u}) dt \quad (9)$$

where $\phi(\mathbf{x}(T_f))$ is the terminal cost, $\psi^j(\mathbf{x}(T_j^-))$ is the via-point cost at the j -th switching instance and $h(\mathbf{x}, \mathbf{u})$ is the running cost.

2.4 Optimal Control of Switching Dynamics and Discrete State Transition

In brief, the iLQR method solves an optimal control problem of the locally linear quadratic approximation of the nonlinear dynamics and the cost function around a nominal trajectory $\bar{\mathbf{x}}$ and control sequence $\bar{\mathbf{u}}$ in discrete time, and iteratively improves the solutions.

In order to incorporate switching dynamics and discrete state transition with a given switching sequence, the hybrid dynamics (1) and (2) are linearised in discrete time around the nominal trajectory and control sequence as

$$\delta \mathbf{x}_{k+1} = \mathbf{A}_k \delta \mathbf{x}_k + \mathbf{B}_k \delta \mathbf{u}_k \quad (10)$$

$$\delta \mathbf{x}_{k_j}^+ = \mathbf{\Gamma}_{k_j} \delta \mathbf{x}_{k_j}^- \quad (11)$$

$$\mathbf{A}_k = \mathbf{I} + \Delta t_j \left. \frac{\partial \mathbf{f}_{i_j}}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}_k}, \quad \mathbf{B}_k = \Delta t_j \left. \frac{\partial \mathbf{f}_{i_j}}{\partial \mathbf{u}} \right|_{\mathbf{u}=\mathbf{u}_k} \quad (12)$$

$$\mathbf{\Gamma}_{k_j} = \left. \frac{\partial \Delta^{i_{j-1}, i_j}}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}_{k_j}^-} \quad (13)$$

where $\delta \mathbf{x}_k = \mathbf{x}_k - \bar{\mathbf{x}}_k$, $\delta \mathbf{u}_k = \mathbf{u}_k - \bar{\mathbf{u}}_k$, k is the discrete time step, Δt_j is the sampling time for the time interval $T_j \leq t < T_{j+1}$, and k_j is the j -th switching instance in the discretised time step.

When using the composite model of the dynamics ($\mathbf{f}_c$) introduced in (8) the linearisation of the dynamics is provided in two parts. The linearisation of $\tilde{\mathbf{f}}$ is obtained by replacing $\mathbf{f}$ with $\tilde{\mathbf{f}}$ in (10) and (12). The derivatives of the learned model ($\tilde{\mathbf{f}}$) are obtained analytically by differentiating with respect to the inputs $\mathbf{z} = (\mathbf{x}; \mathbf{u})$ as suggested in [15]:

$$\frac{\partial \tilde{\mathbf{f}}(\mathbf{z})}{\partial \mathbf{z}} = \frac{1}{W} \sum_k \left(\frac{\partial w_k}{\partial \mathbf{z}} \psi_k(\mathbf{z}) + w_k \frac{\partial \psi_k}{\partial \mathbf{z}} \right) - \frac{1}{W^2} \sum_k w_k(\mathbf{z}) \psi_k(\mathbf{z}) \sum_l \frac{\partial w_l}{\partial \mathbf{z}} \quad (14)$$

$$= \frac{1}{W} \sum_k (-\psi_k w_k \mathbf{D}_k(\mathbf{z} - \mathbf{c}_k) + w_k \mathbf{b}_k) + \frac{\tilde{\mathbf{f}}(\mathbf{z})}{W} \sum_k w_k \mathbf{D}_k(\mathbf{z} - \mathbf{c}_k) \quad (15)$$

where

$$\frac{\partial \tilde{\mathbf{f}}(\mathbf{z})}{\partial \mathbf{z}} = \begin{pmatrix} \partial \tilde{\mathbf{f}} / \partial \mathbf{x} \\ \partial \tilde{\mathbf{f}} / \partial \mathbf{u} \end{pmatrix}. \quad (16)$$

Since there are no changes on the encoding of the instantaneous state transition (2) the equations in (11) and (13) remain unchanged for the iLQG-LD framework.

The composite cost function (9) is locally approximated in a quadratic form as

$$\begin{aligned} \Delta J = & \delta \mathbf{x}_N^T \phi_{\mathbf{x}} + \frac{1}{2} \delta \mathbf{x}_N^T \phi_{\mathbf{xx}} \delta \mathbf{x}_N + \sum_{j=1}^K \left((\delta \mathbf{x}_{k_j}^-)^T \psi_{\mathbf{x}}^j + \frac{1}{2} (\delta \mathbf{x}_{k_j}^-)^T \psi_{\mathbf{xx}}^j \delta \mathbf{x}_{k_j}^- \right) \\ & + \sum_{k=1}^N \left(\delta \mathbf{x}_k^T h_{\mathbf{x}} + \delta \mathbf{u}_k^T h_{\mathbf{u}} + \frac{1}{2} \delta \mathbf{x}_k^T h_{\mathbf{xx}} \delta \mathbf{x}_k + \frac{1}{2} \delta \mathbf{u}_k^T h_{\mathbf{uu}} \delta \mathbf{u}_k + \delta \mathbf{u}_k^T h_{\mathbf{ux}} \delta \mathbf{x}_k \right) \Delta t_j \quad (17) \end{aligned}$$

and a local quadratic approximation of the optimal cost-to-go function is

$$v_k(\delta \mathbf{x}_k) = \frac{1}{2} \delta \mathbf{x}_k^T \mathbf{S}_k \delta \mathbf{x}_k + \delta \mathbf{x}_k^T \mathbf{s}_k. \quad (18)$$

For notational convenience, note that in (17), $\phi_{\mathbf{x}}$ and $\phi_{\mathbf{xx}}$ denote $\phi_{\mathbf{x}} = \frac{\partial \phi}{\partial \mathbf{x}}$ and $\phi_{\mathbf{xx}} = \frac{\partial^2 \phi}{\partial \mathbf{x}^2}$, respectively. Similar definitions apply to other partial derivatives.

The local control law $\delta \mathbf{u}_k$ of the form

$$\delta \mathbf{u}_k = \mathbf{l}_k + \mathbf{L}_k \delta \mathbf{x}_k \quad (19)$$

is obtained from the Bellman equation

$$v_k(\delta \mathbf{x}_k) = \min_{\delta \mathbf{u}} \{h_k(\delta \mathbf{x}_k, \delta \mathbf{u}_k) + v_{k+1}(\delta \mathbf{x}_{k+1})\} \quad (20)$$

by substituting (10) and (18) into the equation (20), where h_k is the local approximation of the running cost in (17) (see [22] for details).

Once we have a locally optimal control command $\delta \mathbf{u}$, the nominal control sequence is updated as $\bar{\mathbf{u}} \leftarrow \bar{\mathbf{u}} + \delta \mathbf{u}$. Then, the new nominal trajectory $\bar{\mathbf{x}}$ is computed by running the obtained control $\bar{\mathbf{u}}$ and the above process is iterated until convergence.

In order to optimise the switching instances and the total movement duration, we introduce a scaling parameter and sampling time for each duration between switching as (cf. (12) and (17)):

$$\Delta t'_j = \frac{1}{\beta_j} \Delta t_j \quad \text{for} \quad T_j \leq t < T_{j+1}, \quad \text{where } j = 0, \dots, K. \quad (21)$$

By optimising the vector of temporal scaling factors $\boldsymbol{\beta} = [\beta_0, \dots, \beta_K]^T$ via gradient descent [23] we obtain each switching instance T_{j+1} and the total movement duration T_f . This approach was applied previously [24, 25] to optimise the frequency of the periodic movement and the movement duration of swing locomotion in a brachiation task.

In the complete optimisation, computation of optimal feedback control law and temporal scaling parameter update are iteratively performed until convergence. A pseudocode of the complete algorithm is summarised in Algorithm 1.

3 Brachiation System Dynamics

We evaluate the effectiveness of the approach on a robot brachiation task which incorporates switching dynamics and multiple phases of the movement in a realistic VSA actuator model. We consider a two-link underactuated brachiating robot with a MACCEPA [26] variable stiffness actuator. The equation of motion of the system used takes the standard form of rigid body dynamics where only the second joint is actuated ⁴:

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) + \mathbf{D}\dot{\mathbf{q}} = \begin{bmatrix} 0 \\ \tau(\mathbf{q}, \mathbf{q}_m) \end{bmatrix} \quad (22)$$

⁴ For notational simplicity, the subscript i is omitted.

Algorithm 1 Complete optimisation algorithm for hybrid dynamics with temporal optimisation

- 1: **Input:**
 - Timed switching plant dynamics $\mathbf{f}_i$ (1 or 8), discrete state transition Δ^{i_{j-1}, i_j} (2) and switching sequence
 - Composite cost function J (9)
 - 2: **Initialise:**
 - Nominal switching instance and final time $T_1, \dots, T_K$ and T_f
 - Nominal control sequence $\bar{\mathbf{u}}$ and corresponding $\bar{\mathbf{x}}$
 - 3: **repeat**
 - 4: **repeat**
 - 5: **Optimise control sequence $\bar{\mathbf{u}}$:**
 - Obtain linearised time-based switching dynamics (10 or 15) and state transition (11) around $\bar{\mathbf{x}}$ and $\bar{\mathbf{u}}$ in discrete time with current Δt_j
 - Compute quadratic approximation of the composite cost (17)
 - Solve local optimal control problem to obtain $\delta \mathbf{u}$ (19)
 - Apply $\delta \mathbf{u}$ to the linearised hybrid dynamics (10) and (17)
 - Update nominal control sequence $\bar{\mathbf{u}} \leftarrow \bar{\mathbf{u}} + \delta \mathbf{u}$, trajectory $\bar{\mathbf{x}}$ and cost J
 - 6: **until** convergence
 - 7: **Temporal optimisation: update Δt_j :**
 - Update the vector of temporal scaling factor β and corresponding sampling time $\Delta t_0, \dots, \Delta t_K$ in (21) via gradient descent [23].
 - 8: **until** convergence
 - 9: **Output:**
 - Optimal feedback control law $\mathbf{u}(\mathbf{x}, t)$: forward optimal control sequence $\mathbf{u}_{opt}$, optimal trajectory $\mathbf{x}_{opt}(t)$ and optimal gain matrix $\mathbf{L}_{opt}(t)$:
 $\mathbf{u}(\mathbf{x}, t) = \mathbf{u}_{opt}(t) + \mathbf{L}_{opt}(t)(\mathbf{x}(t) - \mathbf{x}_{opt}(t))$
 - Optimal switching instance $T_1, \dots, T_K$ and final time T_f
 - Optimal composite cost J
-

where $\mathbf{q} = [q_1, q_2]^T$ is the joint angle vector, $\mathbf{M}$ is the inertia matrix, $\mathbf{C}$ is the Coriolis term, $\mathbf{g}$ is the gravity vector, $\mathbf{D}$ is the viscous damping matrix, τ is the joint torque acting on the second joint given by the VSA, and $\mathbf{q}_m$ is the motor positions vector in the VSA as described below.

The MACCEPA actuator is equipped with two position controlled servo motors, $\mathbf{q}_m = [q_{m1}, q_{m2}]^T$, which control the equilibrium position and the spring pre-tension⁵, respectively. The servo motor dynamics are approximated by a second order system with a PD feedback control, as mentioned in (4):

$$\ddot{\mathbf{q}}_m + 2\alpha\dot{\mathbf{q}}_m + \alpha^2\mathbf{q}_m = \alpha^2\mathbf{u} \quad (23)$$

where $\mathbf{u} = [u_1, u_2]^T$ is the motor position command, α determines the bandwidth of the actuator. In this study, we use $\alpha = \text{diag}(20, 25)$. The range of the commands of the servo motors are limited as $u_1 \in [-\pi/2, \pi/2]$ and $u_2 \in [0, \pi/2]$.

⁵ Which is used to modulate the stiffness of the joint, for details see [26].

We use the model parameters shown in Table 1 and the MACCEPA parameters with the spring constant $\kappa = 771$ N/m, the lever length $B = 0.03$ m, the pin displacement $C = 0.125$ m and the drum radius $r_d = 0.01$ m.

Fig. 3. Two-link brachiating robot model with the VSA joint with the inertial and geometric parameters. The parameters of the robot are given in Table 1, where the indices i denote the link number in this figure and Table 1.

Robot parameters		i=1	i=2	$i=1$
Mass	m_i (kg)	1.390	0.527	<i>1.240</i>
Moment of inertia	I_i (kgm ²)	0.0297	0.0104	<i>0.0278</i>
Link length	l_i (m)	0.46	0.46	<i>0.46</i>
COM location	l_{ci} (m)	0.362	0.233	<i>0.350</i>
Viscous friction	d_i (Nm/s)	0.03	0.035	<i>0.03</i>

Table 1. Model parameters of the two-link brachiating robot. The final column shows the change of parameters of the first link of the system under the changed mass distribution described in Sec. 4

4 Experimental setup

To test the efficiency of our approach we create a scenario where the difference between the true and the assumed model is caused by a change in the mass (and implicitly mass distribution) on one the links (i.e. the mass of the true model is smaller by 150 g (located at the joint) on link $i = 1$). The changed model parameters are shown in the right column of Table 1 ⁶.

Due to the nature of the discrepancy introduced, the error in the dynamics manifests itself only in the joint accelerations. Thus, we require to map the error

⁶ The MACCEPA parameters are the same as described in the previous section.

just in those two dimensions, reducing the dimension of the output of the LWPR model $\tilde{\mathbf{f}}$ from $n = 8$ to 2, where the predictions are added on the corresponding dimension of $\tilde{\mathbf{f}}$ within $\mathbf{f}_c$ (8). Note that different discrepancies will necessitate estimation of the full 8-dim state error.

In line with previous work, we will demonstrate the effectiveness of the proposed approach on a multi-phase, asymmetric swing-up and brachiation task with a VSA while incorporating *continuous, online model learning*. Specifically, in the multi-phase task, the robot swings up from the suspended posture to the target at $d_1 = 0.40$ m and subsequently moves to the target located at $d_2 = 0.42$ m and $d_3 = 0.46$ m, respectively.

Since the system has an asymmetric configuration and the state space of the swing up task is significantly different from that of a brachiation movement we proceed by first learning a separate error model for each phase. The procedure used is briefly described in Algorithm 2. The initial exploration loop is performed in order to pre-train the LWPR model $\tilde{\mathbf{f}}_i$ (as an alternative to applying motor babbling), the later loop is using iLQG-LD to refine the model in an online fashion. In our experiments the training data is obtained by using a simulated version of the true dynamics, which is an analytic model incorporating the discrepancy.

Algorithm 2 Description of the learning and exploration procedure

Given:

- analytic dynamics for one configuration $\tilde{\mathbf{f}}_i$ and start state $\mathbf{x}_0$
- thresholds for target reaching ε_T and model accuracy ε_M
- the associated cost function J (including desired target $\mathbf{x}_T$)
- p number of initial exploration training episodes

Initialise

- $\tilde{\mathbf{f}}_i(\mathbf{x}, \mathbf{u})$; $\mathbf{f}_{c_i}(\mathbf{x}, \mathbf{u}) = \tilde{\mathbf{f}}_i(\mathbf{x}, \mathbf{u}) + \tilde{\mathbf{f}}_i(\mathbf{x}, \mathbf{u})$

repeat

generate $\bar{\mathbf{u}}, \bar{\mathbf{x}}, \mathbf{L}$ using $\tilde{\mathbf{f}}_i(\mathbf{x}, \mathbf{u})$ for an artificial target (a new target at each iteration obtained by sampling around $\mathbf{x}_T$)

apply the solution to the true dynamics and train the model on the collected data

until p training episodes have been performed

repeat

apply iLQG-LD for target $\mathbf{x}_T$

until ε_T and ε_M conditions are met

4.1 Individual phase learning

Using the traditional OC framework in the presence of an accurate dynamics model, the multi-phase task described previously was achieved with a position error of just 0.002 m. Once the discrepancy detailed in Sec. 4 is introduced, the planned solution is no longer valid and the final position deviates from the desired target (Fig. 5, blue line). We deploy the iLQG-LD framework in order to learn the new behaviour of the system and recover the task performance.

As a measure of the model accuracy we use the *normalised mean square error* (nMSE) of the model prediction on the true optimal trajectory (if given access to the analytic form of the true dynamics of the system). The nMSE is defined as $nMSE(y, \tilde{y}) = \frac{1}{n\sigma_y^2} \sum_{i=1}^n (y_i - \tilde{y}_i)^2$ where y is the desired output data set of size n and $\tilde{y}$ represents the LWPR predictions. The evolution of the nMSE at each stage of the training for every phase is shown in Fig. 4.

Fig. 4. Evolution of the nMSE for each phase of the movement, at each episode

In the first part (pre-training phase in Fig. 4) we generate random targets around the desired x_T . A movement is planned for these targets using the assumed model ($\tilde{f}$). The obtained command solution is then applied to the simulated version of the true dynamics, using a closed loop control scheme. We repeat the procedure for a set of 10 noise contaminated versions of the commands. The collected data is used to train the model.

This pre-training phase seeds the model with information within the region of interest, prior to using it for planning. This reduces the load on the iLQG-LD by lowering the number of iterations required for convergence. For each phase of the movement, at the end of the procedure, the planned trajectory matched the

behaviour obtained from running the command solution on the real plant (the final $nMSE$ has an order of magnitude of 10^{-4}).

Overall the discrepancy is small enough to allow reaching the desired end effector position within a threshold of $\varepsilon_T = 0.040$ m accuracy⁷. Fig. 5 shows the effect of the learning by comparing the performance of the planning with the erroneous model and with the composite model obtained after training.

Fig. 5. Comparison of the final position achieved (for each individual phase) when using the initial planning (erroneous model - blue, thick line) and the final planning (composite model - black, thin line). Intermediary solutions obtained at each step of the iLQG-LD run are depicted in grey.

4.2 Multi-phase performance

In the previous section we showed that our approach to iLQG-LD is able to cope with the requirements of the task in each phase of the movement. For a full solution we use the newly learned models from each phase to obtain the global solution for the multi-phase task wrt. the composite cost function J (9). We use the phase optimal solutions obtained at the previous stage as the initial command sequence, the resulting behaviour is displayed in Fig. 6. The planner is able to use the learned model to achieve the intermediary and final goals, while the expected behaviour provides a reliable match to the actual system's behaviour⁸. The cost of multi-phase optimised solution ($J = 39.17$) is significantly lower than the sum of the costs of the individual phase solutions ($J = 58.23$).

⁷ The error in the swing up task is 0.033 m, while for brachiations the value is 0.004 m. In future work we aim to bring the former value to the same magnitude.

⁸ We consider that if the position at the end of each phase is within our prescribed threshold $\varepsilon_T = 0.040$ of the desired target the system is able to start the next phase from the ideal location, thus resembling the effect of the grabber on the hardware.

Fig. 6. Performance of the fully optimised multi-phase task using the composite model. Thick grey lines: planned movement. Black and blue thin lines: actual system movement.

4.3 Performance of learning

In the previous experiment, we investigated a single (arbitrarily chosen) mass distribution discrepancy. Next we investigate the capacity of our approach to cope with a wider range of mismatched dynamics. For this, we consider the magnitude of the change that is bounded by the capability of the altered system to achieve all the phases of the movement presented in Fig. 6. We define these bounds as the limit values of mass change that allow the same accuracy in task execution, under the same cost function (9). The corresponding values for these limits, found empirically, are -0.200 kg and $+0.300$ kg, respectively ⁹.

We apply our framework to the scenarios where the mass has been altered to these boundary values and demonstrate the result on just one of the phases, namely the first brachiation move, to study the relative effects.

Fig. 7.a shows the effect of the alteration introduced, when executing the commands resulting from the initial planning (using the erroneous model). After training and re-planning the model starts approximating the true behaviour of the system, such that in less than 10 episodes, the system is once again able to reach the desired target, as depicted in Fig. 7.b and 7.c.

⁹ We note that in our experiments, we assume that modulating the mass distribution does not affect the motor dynamics – this represents a simplified scenario. In the real hardware, the speed of the motor dynamics (4) is indeed a function of the overall load distribution.

Fig. 7. (a): Effect of the discrepancy introduced by the limit values on the mass modulation (blue lines). The behaviour when the model match is correct is depicted for comparison (thin green line). **(b) - (c):** Comparison of the final position achieved (for each individual phase) when using the initial planning (erroneous model - thick blue line) and the final planning (composite model - thin black line). Intermediary solutions obtained at each step of the iLQG-LD run are depicted in grey. Results for the boundary value discrepancies on the mass distribution.

5 Conclusion

In this work we have presented an extension of our methodology for movement optimisation with multiple phases and switching dynamics, including variable impedance actuators. We broaden the approach by incorporating adaptive learning, which allows for adjustments to the dynamics model, based on changes occurred to the system's behaviour, or when the behaviour cannot be fully capture by a rigid body dynamics formulation. In future work we aim to investigate a wider range of model discrepancies and show the performance of the extended approach on a hardware implementation.

References

- [1] B. Siciliano and O. Khatib, *Springer Handbook of Robotics*. Springer, 2008.
- [2] O. Sigaud, C. Salaün, and V. Padois, "On-line regression algorithms for learning mechanical models of robots: a survey," *Rob. and Auton. Sys.*, vol. 59, 2011.
- [3] L. Sciavicco and L. Villani, *Robotics: modelling, planning & control*. Springer, 2009.
- [4] C. G. Atkeson, C. H. An, and J. M. Hollerbach, "Estimation of inertial parameters of manipulator loads and links," *The Int. Jour. of Rob. Research*, vol. 5, 1986.
- [5] W. Khalil and É. Dombre, *Modeling, identification & control of robots*. Kogan Page Science paper edition, Kogan Page Science, 2004.
- [6] J. Nakanishi, J. A. Farrell, and S. Schaal, "Composite adaptive control with locally weighted statistical learning," *Neural Networks*, vol. 18, no. 1, pp. 71–90, 2005.

- [7] D. Mitrovic, S. Klanke, M. Howard, and S. Vijayakumar, "Exploiting sensorimotor stochasticity for learning control of variable impedance actuators," in *Humanoids*, 2010.
- [8] D. Nguyen-Tuong and J. Peters, "Model learning for robot control: a survey," *Cognitive processing*, vol. 12, no. 4, pp. 319–340, 2011.
- [9] S. Vijayakumar, A. D'souza, T. Shibata, J. Conradt, and S. Schaal, "Statistical learning for humanoid robots," *Auton. Rob.*, vol. 12, no. 1, pp. 55–69, 2002.
- [10] D. Nguyen-Tuong, M. Seeger, and J. Peters, "Model learning with local gaussian process regression," *Advanced Robotics*, vol. 23, no. 15, pp. 2015–2034, 2009.
- [11] J.-A. Ting, M. Kalakrishnan, S. Vijayakumar, and S. Schaal, "Bayesian kernel shaping for learning control," in *Advances in NIPS*, pp. 1673–1680, 2009.
- [12] O. Herbort, M. V. Butz, and G. Pedersen, "The sure_reach model for motor learning and control of a redundant arm: from modeling human behavior to applications in robotics," in *From motor learning to interaction learning in robots*, pp. 85–106, Springer, 2010.
- [13] S. Schaal, C. G. Atkeson, and S. Vijayakumar, "Scalable techniques from non-parametric statistics for real time robot learning," *Applied Intelligence*, vol. 17, no. 1, pp. 49–60, 2002.
- [14] M. Kalakrishnan, J. Buchli, P. Pastor, and S. Schaal, "Learning locomotion over rough terrain using terrain templates," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS 2009)*, 2009.
- [15] C. G. Atkeson, A. W. Moore, and S. Schaal, "Locally weighted learning for control," in *Artificial Intelligence Review*, vol. 11, pp. 75–113, Springer, 1997.
- [16] D. Mitrovic, S. Klanke, and S. Vijayakumar, "Adaptive optimal feedback control with learned internal dynamics models," *From Motor Learning to Interaction Learning in Robots*, pp. 65–84, 2010.
- [17] D. Mitrovic, S. Klanke, and S. Vijayakumar, "Optimal control with adaptive internal dynamics models," 2008.
- [18] S. Klanke, S. Vijayakumar, and S. Schaal, "A library for locally weighted projection regression," *The Jour. of Machine Learning Research*, vol. 9, 2008.
- [19] J. Nakanishi, A. Radulescu, and S. Vijayakumar, "Spatio-temporal optimization of multi-phase movements: Dealing with contacts and switching dynamics," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS 2013)*, 2013.
- [20] T. M. Caldwell and T. D. Murphey, "Switching mode generation and optimal estimation with application to skid-steering," *Automatica*, vol. 47, 2011.
- [21] D. Braun, M. Howard, and S. Vijayakumar, "Optimal variable stiffness control: formulation and application to explosive movement tasks," *Auton. Rob.*, 2012.
- [22] W. Li and E. Todorov, "Iterative linearization methods for approximately optimal control and estimation of non-linear stochastic system," *International Journal of Control*, vol. 80, no. 9, pp. 1439–1453, 2007.
- [23] K. Rawlik, M. Toussaint, and S. Vijayakumar, "An approximate inference approach to temporal optimization in optimal control," in *Advances in NIPS 23*, pp. 2011–2019, 2010.
- [24] J. Nakanishi, K. Rawlik, and S. Vijayakumar, "Stiffness and temporal optimization in periodic movements: An optimal control approach," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 718–724, 2011.
- [25] J. Nakanishi and S. Vijayakumar, "Exploiting passive dynamics with variable stiffness actuation in robot brachiation," in *Robotics: Science and Systems*, 2012.
- [26] R. Van Ham *et al.*, "MACCEPA: Design and implementation in a biped robot," *Rob. and Auton. Sys.*, vol. 55 (10), pp. 761–768, 2007.